

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

- 1 1. (Currently amended) A method for selecting a node to host a primary
2 server for a service from a plurality of nodes in a distributed computing system,
3 the method comprising:
 - 4 periodically sending checkpoint information from a primary server to a
5 secondary server;
 - 6 receiving an indication that a state of the distributed computing system has
7 changed;
 - 8 in response to the indication, determining if there is already a node hosting
9 the primary server for the service; and
 - 10 if there is not already a node hosting the primary server, selecting a node to
11 host the primary server based upon rank information for the nodes, wherein the
12 rank information specifies whether the selected node is a secondary server which
13 has received the checkpoint information, whereby the secondary server is able to
14 take over for the primary server without having to wait to receive additional
15 configuration information;
 - 16 wherein selecting a node to host the primary server based upon rank
17 information for the nodes occurs concurrently on all active nodes; and
18 wherein the identity of the node that hosts the primary server is
19 communicated through a set of lockable, shared candidate variables located on the
20 individual nodes that are accessed remotely by all of the active nodes to disqualify
21 individual nodes from hosting the primary server based upon rank information.

1 2. (Original) The method of claim 1, wherein selecting the node to host the
2 primary server involves:

3 assuming that a given node from the plurality of nodes in the distributed
4 computing system hosts the primary server,

5 communicating rank information between the given node and other nodes
6 in the distributed computing system, wherein each node in the distributed
7 computing system has a unique rank with respect to the other nodes in the
8 distributed computing system,

9 comparing a rank of the given node with a rank of the other nodes in the
10 distributed computing system, and

11 if one of the other nodes in the distributed computing system has a higher
12 rank than the given node, disqualifying the given node from hosting the primary
13 server.

1 3. (Original) The method of claim 2, further comprising, if there exists a
2 node that is configured to host the primary server, allowing the node that is
3 configured to host the primary server to communicate with other nodes in the
4 distributed computing system in order to disqualify the other nodes from hosting
5 the primary server.

1 4. (Original) The method of claim 2, wherein assuming that the given node
2 hosts the primary server involves:

3 maintaining a candidate variable in the given node identifying a candidate
4 node to host the primary server; and

5 initially setting the candidate variable to identify the given node.

1 5. (Original) The method of claim 1, further comprising, after a new node
2 has been selected to host the primary server, if the new node is different from a

3 previous node that hosted the primary server, establishing connections for the
4 service to the new node.

1 6. (Original) The method of claim 1, further comprising, after a new node
2 has been selected to host the primary server, if the new node is different from a
3 previous node that hosted the primary server, configuring the new node to host the
4 primary server for the service.

1 7. (Original) The method of claim 1, further comprising restarting the
2 service if the service was interrupted as a result of the change in state of the
3 distributed computing system.

1 8. (Original) The method of claim 2, wherein the given node in the
2 distributed computing system can be dynamically configured to act as one of:
3 a host for the primary server for the service;
4 a host for a secondary server for the service, wherein the secondary server
5 periodically receives checkpointing information from the primary server; and
6 a spare for the primary server, wherein the spare does not receive
7 checkpointing information from the primary server.

1 9. (Original) The method of claim 8, further comprising, upon initial
2 startup of the service, selecting a highest ranking spare to host the primary server
3 for the service.

1 10. (Original) The method of claim 8, further comprising allowing the
2 primary server to configure spares in the distributed computing system to host
3 secondary servers for the service.

1 11. (Original) The method of claim 8, wherein comparing the rank of the
2 given node with the rank of the other nodes in the distributed computing system
3 involves considering a host for the primary server to have a higher rank than a
4 host for a space, and considering a host for a secondary server to have a higher
5 rank than a spare.

1 12. (Original) The method of claim 2, wherein disqualifying the given
2 node from hosting the primary server involves ceasing to communicate rank
3 information between the given node and the other nodes in the distributed
4 computing system.

1 13. (Currently amended) A computer-readable storage medium storing
2 instructions that when executed by a computer cause the computer to perform a
3 method for selecting a node to host a primary server for a service from a plurality
4 of nodes in a distributed computing system, the method comprising:

5 periodically sending checkpoint information from a primary server to a
6 secondary server;

7 receiving an indication that a state of the distributed computing system has
8 changed;

9 in response to the indication, determining if there is already a node hosting
10 the primary server for the service; and

11 if there is not already a node hosting the primary server, selecting a node to
12 host the primary server based upon rank information for the nodes, wherein the
13 rank information specifies whether the selected node is a secondary server which
14 has received the checkpoint information, whereby the secondary server is able to
15 take over for the primary server without having to wait to receive additional
16 configuration information;

17 wherein selecting a node to host the primary server based upon rank
18 information for the nodes occurs concurrently on all active nodes; and
19 wherein the identity of the node that hosts the primary server is
20 communicated through a set of lockable, shared candidate variables located on the
21 individual nodes that are accessed remotely by all of the active nodes to disqualify
22 individual nodes from hosting the primary server based upon rank information.

1 14. (Original) The computer-readable storage medium of claim 13,
2 wherein selecting the node to host the primary server involves:
3 assuming that a given node from the plurality of nodes in the distributed
4 computing system hosts the primary server,
5 communicating rank information between the given node and other nodes
6 in the distributed computing system, wherein each node in the distributed
7 computing system has a unique rank with respect to the other nodes in the
8 distributed computing system,
9 comparing a rank of the given node with a rank of the other nodes in the
10 distributed computing system, and
11 if one of the other nodes in the distributed computing system has a higher
12 rank than the given node, disqualifying the given node from hosting the primary
13 server.

1 15. (Original) The computer-readable storage medium of claim 14,
2 wherein if there exists a node that is configured to host the primary server, the
3 method further comprises allowing the node that is configured to host the primary
4 server to communicate with other nodes in the distributed computing system in
5 order to disqualify the other nodes from hosting the primary server.

1 16. (Original) The computer-readable storage medium of claim 14,
2 wherein assuming that the given node hosts the primary server involves:
3 maintaining a candidate variable in the given node identifying a candidate
4 node to host the primary server; and
5 initially setting the candidate variable to identify the given node.

1 17. (Original) The computer-readable storage medium of claim 13,
2 wherein after a new node has been selected to host the primary server, if the new
3 node is different from a previous node that hosted the primary server, the method
4 further comprises establishing connections for the service to the new node.

1 18. (Original) The computer-readable storage medium of claim 13,
2 wherein after a new node has been selected to host the primary server, if the new
3 node is different from a previous node that hosted the primary server, the method
4 further comprises configuring the new node to host the primary server for the
5 service.

1 19. (Original) The computer-readable storage medium of claim 13,
2 wherein the method further comprises restarting the service if the service was
3 interrupted as a result of the change in state of the distributed computing system.

1 20. (Previously amended) The computer-readable storage medium of claim
2 14, wherein the given node in the distributed computing system can be
3 dynamically configured to act as one of:
4 a host for the primary server for the service;
5 a host for a secondary server for the service, wherein the secondary server
6 periodically receives checkpointing information from the primary server; and

7 a spare for the primary server, wherein the spare does not receive
8 checkpointing information from the primary server.

1 21. (Original) The computer-readable storage medium of claim 20,
2 wherein upon initial startup of the service, the method further comprises selecting
3 a highest ranking spare to host the primary server for the service.

1 22. (Original) The computer-readable storage medium of claim 20,
2 wherein the method further comprises allowing the primary server to configure
3 spares in the distributed computing system to host secondary servers for the
4 service.

1 23. (Original) The computer-readable storage medium of claim 20,
2 wherein comparing the rank of the given node with the rank of the other nodes in
3 the distributed computing system involves considering a host for the primary
4 server to have a higher rank than a host for a space, and considering a host for a
5 secondary server to have a higher rank than a spare.

1 24. (Original) The computer-readable storage medium of claim 14,
2 wherein disqualifying the given node from hosting the primary server involves
3 ceasing to communicate rank information between the given node and the other
4 nodes in the distributed computing system.

1 25. (Currently amended) An apparatus that selects a node to host a primary
2 server for a service from a plurality of nodes in a distributed computing system,
3 the apparatus comprising:
4 a sending mechanism configured to periodically send checkpoint
5 information from a primary server to a secondary server;

6 a receiving mechanism that is configured to receive an indication that a
7 state of the distributed computing system has changed;
8 a determination mechanism that is configured to determine if there is
9 already a node hosting the primary server for the service in response to the
10 indication;
11 a selecting mechanism, wherein if there is not already a node hosting the
12 primary server, the selecting mechanism is configured to select a node to host the
13 primary server based upon rank information for the nodes, wherein the rank
14 information specifies whether the selected node is a secondary server which has
15 received the checkpoint information, whereby the secondary server is able to take
16 over for the primary server without having to wait to receive additional
17 configuration information;
18 wherein selecting a node to host the primary server based upon rank
19 information for the nodes occurs concurrently on all active nodes; and
20 wherein the identity of the node that hosts the primary server is
21 communicated through a set of lockable, shared candidate variables located on the
22 individual nodes that are accessed remotely by all of the active nodes to disqualify
23 individual nodes from hosting the primary server based upon rank information.

1 26. (Original) The apparatus of claim 25, wherein, in selecting a node to
2 host the primary server based upon rank information, the selecting mechanism is
3 configured to:
4 communicate rank information between the given node and other nodes in
5 the distributed computing system, wherein each node in the distributed computing
6 system has a unique rank with respect to the other nodes in the distributed
7 computing system, and to
8 compare a rank of the given node with a rank of the other nodes in the
9 distributed computing system.

1 27. (Original) The apparatus of claim 26, further comprising a
2 disqualification mechanism that is configured to disqualify the given node from
3 hosting the primary server if one of the other nodes in the distributed computing
4 system has a higher rank than the given node.

1 28. (Original) The apparatus of claim 26, further comprising a mechanism
2 on the primary server that is configured to communicate with other nodes in the
3 distributed computing system in order to disqualify the other nodes from hosting
4 the primary server.

1 29. (Original) The apparatus of claim 26, wherein the selecting mechanism
2 is configured to:
3 maintain a candidate variable in the given node identifying a candidate
4 node to host the primary server; and to
5 initially set the candidate variable to identify the given node.

1 30. (Original) The apparatus of claim 25, further comprising a connection
2 mechanism that is configured to establish connections for the service to a new
3 node after the new node has been selected to host the primary server, and if the
4 new node is different from a previous node that hosted the primary server.

1 31. (Original) The apparatus of claim 25, further comprising a mechanism
2 that configures a new node to host the primary server for the service, after the new
3 node has been selected to host the primary server, and if the new node is different
4 from a previous node that hosted the primary server.

1 32. (Original) The apparatus of claim 25, further comprising a restarting
2 mechanism that is configured to restart the service if the service was interrupted as
3 a result of the change in state of the distributed computing system.

1 33. (Previously amended) The apparatus of claim 26, wherein the given
2 node in the distributed computing system can be dynamically configured to act as
3 one of:

4 a host for the primary server for the service;
5 a host for a secondary server for the service, wherein the secondary server
6 periodically receives checkpointing information from the primary server; and
7 a spare for the primary server, wherein the spare does not receive
8 checkpointing information from the primary server.

1 34. (Original) The apparatus of claim 33, further comprising an
2 initialization mechanism wherein during initialization of the service, the
3 initialization mechanism is configured to select a highest ranking spare to host the
4 primary server for the service.

1 35. (Original) The apparatus of claim 33, further comprising a promotion
2 mechanism on the primary server that is configured to promote spares in the
3 distributed computing system to host secondary servers for the service.

1 36. (Original) The apparatus of claim 33, wherein while comparing the
2 rank of the given node with the rank of the other nodes in the distributed
3 computing system, the selecting mechanism is configured to consider a host for
4 the primary server to have a higher rank than a host for a secondary server, and to
5 consider a host for a secondary server to have a higher rank than a spare.

1 37. (Original) The apparatus of claim 26, wherein the selecting mechanism
2 is configured to cease to communicate rank information between the given node
3 and the other nodes in the distributed computing system after the given node is
4 disqualified by the disqualification mechanism.

1 38. (Currently amended) A method for selecting a node to host a primary
2 server for a service from a plurality of nodes in a distributed computer system,
3 comprising:

4 periodically sending checkpoint information from a primary server to a
5 secondary server;

6 communicating disqualification information between the node and
7 remaining nodes in the plurality of nodes through a set of lockable, shared
8 candidate variables located on the individual nodes that are accessed remotely by
9 all of the active nodes;

10 disqualifying the node from hosting the primary server based upon the
11 disqualification information received from the remaining nodes;

12 wherein communicating disqualification information between the node
13 and remaining nodes in the plurality of nodes occurs concurrently on all active
14 nodes;

15 wherein disqualifying the node from hosting the primary server based
16 upon the disqualification information received from the remaining nodes occurs
17 concurrently on all active nodes.

1 39. (Original) The method of claim 38, wherein the disqualification
2 information comprises a node rank information.

1 40. (Original) The method of claim 39, wherein the node rank for a given
2 node is calculated using an assumption that the given node hosts the primary
3 server.

1 41. (Original) The method of claim 40, wherein the calculated node rank is
2 unique with respect to the ranks of other nodes in the distributed computer system.

1 42. (Original) The method of claim 39, wherein the disqualifying of the
2 node comprises:

3 comparing a rank of the node to a set of ranks of the remaining nodes in
4 the distributed computer system; and

5 disqualifying the node from hosting the primary server if one of the set of
6 ranks of the remaining nodes is higher than the rank of the node.

1 43. (Original) The method of claim 38, further comprising repeating the
2 acts of communicating disqualification information and disqualifying the node for
3 at least one more node in the plurality of nodes.